arXiv:1304.5257v2 [cs.SE] 26 Apr 2013

What Makes Code Hard to Understand?

Michael Hansen (mihansen@indiana.edu)
School of Informatics and Computing, 2719 E. 10th Street
Bloomington, IN 47408 USA

Robert L. Goldstone (rgoldsto@indiana.edu)
Dept. of Psychological and Brain Sciences, 1101 E. 10th Street
Bloomington, IN 47405 USA

Andrew Lumsdaine (lums@indiana.edu)
School of Informatics and Computing, 2719 E. 10th Street
Bloomington, IN 47408 USA

April 29, 2013

Abstract

What factors impact the comprehensibility of code? Previous research suggests that expectation-congruent programs should take less time to understand and be less prone to errors. We present an experiment in which participants with programming experience predict the exact output of ten small Python programs. We use subtle differences between program versions to demonstrate that seemingly insignificant notational changes can have profound effects on correctness and response times. Our results show that experience increases performance in most cases, but may hurt performance significantly when underlying assumptions about related code statements are violated.

Keywords: program comprehension; psychology of programming; code complexity.

1 Introduction

The design, creation and interpretation of computer programs are some of the most cognitively challenging tasks that humans perform. Understanding the factors that impact the cognitive complexity of code is important for both applied and theoretical reasoning. Practically, an enormous amount of time is spent developing programs, and even more time is spent debugging them, and so if we can identify factors that expedite these activities, a large amount of time and money can be saved. Theoretically, programming is an excellent task for studying representation, working memory, planning, and problem solving in the real world.

We present a web-based experiment in which participants with a wide variety of Python and overall programming experience predict the output of ten small Python programs. Most of the program texts are less than 20 lines long and have fewer than 8 linearly independent paths (known as cyclomatic complexity [7]). Each program *type* has two or three *versions* with subtle differences that do not significantly change their lines of code (LOC) or cyclomatic complexities (CC). For each participant and program, we grade text responses on a 10-point scale, and record the amount of time taken. The different versions of our programs were designed to test a couple of underlying questions. First, "How are programmers affected by programs that violate their expectations, and does this vary with expertise?" Previous research suggests that programs that violate expectations should take longer to process and be more error-prone

than expectation-congruent programs. There are reasons to expect this benefit for expectation-congruency to interact with experience in opposing ways. Experienced programmers may show a larger influence of expectations due to prolonged training, but they may also have more untapped cognitive resources available for monitoring expectation violations. In fact, given the large percentage of programming time that involves debugging (it is a common saying that 90% of development time is spent debugging 10% of the code), experienced programmers may have developed dedicated monitors for certain kinds of expectation-violating code.

The second question is: "How are programmers influenced by physical characteristics of notation, and does this vary with expertise?" Programmers often feel like the physical properties of notation have only a minor influence on their interpretation process. When in a hurry, they frequently dispense with recommended variable naming, indentation, and formatting as superficial and inconsequential. However, in other formal reasoning domains such as math [5], apparently superficial formatting influences such as physical spacing between operators has been shown to have a profound impact on performance. Furthermore, there is an open question as to whether experienced or inexperienced programmers are more influenced by these physical aspects of code notation. Experienced programmers may show less influence of these "superficial" aspects because they are responding to the deep structure of the code. By contrast, in math reasoning, experienced individuals sometimes show more influence of notational properties of the symbols, apparently because they use perception-action shortcuts involving these properties in order to attain efficiency [5].

2 Related Work

Psychologists have been studying programmers for at least forty years. Early research focused on correlations between task performance and human/language factors, such as how the presence of code comments impacts scores on a program comprehension questionnaire. More recent research has revolved around the cognitive processes underlying program comprehension. Effects of expertise, task, and available tools on program understanding have been found [4]. Studies with experienced programmers have revealed conventions, or "rules of discourse," that can have a profound impact (sometimes negative) on expert program comprehension [8].

Our present research focuses on programs much less complicated than those the average professional programmer typically encounters on a daily basis. The demands of our task are still high, however, because participants must predict precise program output. In this way, it is similar to debugging a short snippet of a larger program. Code studies often take the form of a code review, where programmers must locate errors or answer comprehension questions after the fact (e.g., does the program define a Professor class? [1]). Our task differs by asking programmers to mentally simulate code without necessarily understanding its purpose. In most programs, we intentionally use meaningless identifier names where appropriate (variables `a`, `b`, etc.) to avoid influencing the programmer's mental model.

Similar research has asked beginning (CS1) programming students to read and write code with simple goals, such as the Rainfall Problem [6]. To solve it, students must write a program that averages a list of numbers (rainfall amounts), where the list is terminated with a specific value – e.g., a negative number or 999999. CS1 students perform poorly on the Rainfall Problem across institutions around the world, inspiring researchers to seek better teaching methods. Our work includes many Python novices with a year or less of experience (94 out of 162), so our results may contribute to ongoing research in early programming education.

3 Methods

One hundred and sixy-two participants (129 males, 30 females, 3 unreported) were recruited from the Bloomington, IN area (29), on Amazon's Mechanical Turk (130), and via e-mail (3). All participants were required to have some experience with Python, though we welcomed beginners. The mean age was 28.4 years, with an average of 2.0 years of self-reported Python experience and 6.9 years of programming experience overall. Most of the participants had a college degree (69.8%), and were current or former Computer Science majors (52.5%). Participants from Bloomington were paid $10, and performed the experiment in front of an eye-tracker (see Future Work). Mechanical Turk participants were paid $0.75.

The experiment consisted of a pre-test survey, ten trials (one program each), and a post-test survey. Before the experiment began, participants were given access to a small Python "refresher," which listed the code and output of several small programs. The pre-test survey gathered information about the participant's age, gender, education, and experience. Participants were then asked to predict the printed output of ten short Python programs, one version randomly chosen from each of ten program types (Figure 1). The presentation order and names of the programs were randomized, and all answers were final. Although every program produced error-free output, participants were not informed of this fact beforehand. The post-test survey gauged a participant's confidence in their answers and the perceived difficulty of the task.

We collected a total of 1,602 trials from 162 participants starting November 20, 2012 and ending January 19, 2013. Trials were graded semi-automatically using a custom grading program. A grade of 10 points was assigned for responses that exactly matched the program's output (1,007 out of 1,602 trials). A *correct* grade of 7-9 points was given when responses had the right numbers or letters, but incorrect formatting – e.g., wrong whitespace, commas, brackets. Common errors were given partial credit from 2 to 4 points, depending on correct formatting. All other responses were manually graded by two graders whose instructions were to give fewer than 5 points for incorrect responses, and to take off points for incorrect formatting (clear intermediary calculations or comments were ignored). Graders' responses were strongly correlated ($r(598) = 0.90$), so individual trial grades were averaged. Trial response times ranged from 14 to 256 seconds. Outliers beyond two standard deviations of the mean (in log space) were discarded (60 of 1,602 trials). Participants had a total of 45 minutes to complete the entire experiment (10 trials + surveys), and were required to give an answer to each question.

We had a total of twenty-five Python programs in ten different categories. These programs were designed to be understandable to a wide audience, and therefore did not touch on Python features outside of a first or second introductory programming course. The programs ranged in size from 3 to 24 lines of code (LOC). Their cyclomatic complexities (CC) ranged from 1 to 7, and were moderately correlated with LOC ($r(25) = 0.46$, $p < 0.05$). CC was computed using the open source package PyMetrics.

Mechanical Turk One hundred and thirty participants from Mechanical Turk completed the experiment. Workers were required to pass a Python pre-test, and could only participate once. All code was displayed as an image, making it difficult to copy/paste the code into a Python interpreter for quick answers. All responses were manually screened, and restarted trials or unfinished experiments were discarded.

4 Results

Data analysis was performed in R, and all regressions were done with R's built-in `lm` and `glm` functions. For linear regressions involving the dependent measures grade and RT, we report

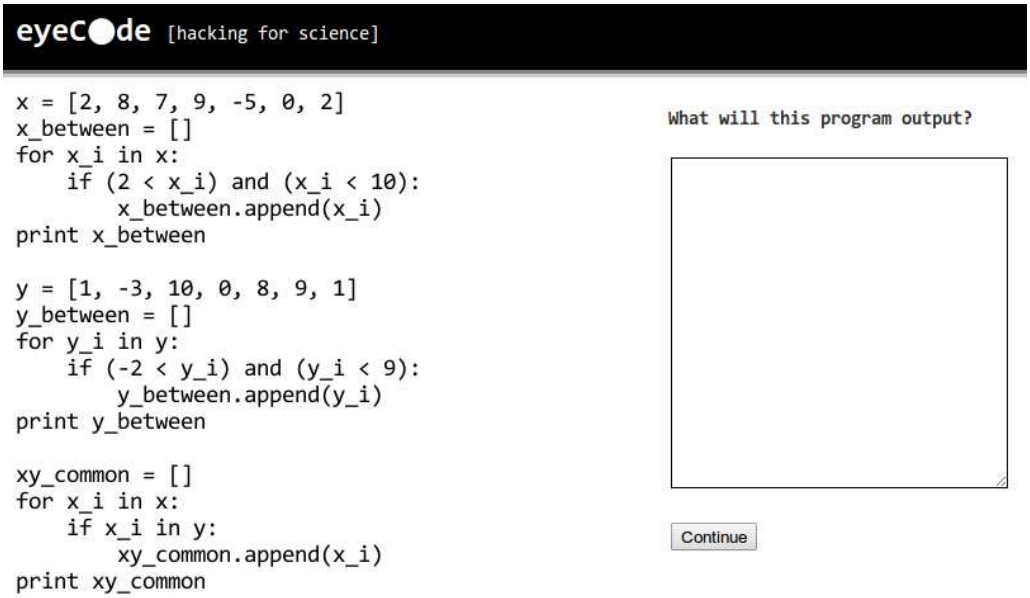


Figure 1: Sample trial from the experiment (`between inline`). Participants were asked to predict the exact output of ten Python programs.

intercepts and coefficients (β). For logistic regressions (probability of a correct answer or a common error), we report base predictor levels and the odds ratios (OR). Odds ratios can be difficult to interpret, and are often confused with relative risk [3]. While the direction of an effect is usually apparent, we caution against their interpretation as effect sizes (especially when OR < 1).

Table 1 summarizes the results in terms of average grades, mean response times (RT), and significant effects (discussed in detail below). Participants did well overall, and the plurality of trials resulted in perfect responses (1,007 out of 1,602). Years of Python experience was a significant predictor of overall grade (intercept = 77.0, $\beta = 1.23$, $p < 0.05$), and a highly significant predictor of giving a perfect response (base = 1.42, OR = 1.09, $p < 0.01$). We discuss grade and RT differences between program versions below.

between (2 versions) This program filters two lists of integers (`x` and `y`), prints the results, and then prints the numbers that `x` and `y` have in common. The `functions` version abstracts the `between` and `common` operations into reusable functions, while the `inline` version inlines the code, duplicating as necessary.

Because this is the longest and most complex program (in terms of CC), we expected more experienced programmers to be faster and to make fewer errors. We were surprised to find a significant effect of Python experience on the probability of making a very specific error (base = 0.13, OR = 1.44, $p < 0.5$). Instead of `[8, 9, 0]` for their last line of output, 22% of participants wrote `[8]` (for both versions). After the experiment, one participant reported they mistakenly performed the "common" operation on lists `x_btwn` and `y_btwn` instead of `x` and `y` because it seemed like the next logical step. If others made the same mistake, this may suggest an addition to Soloway's Rules of Discourse [8]: later computations should follow from earlier ones. We hypothesize that moving the `common` operation code before the two instances of `between` would eliminate this error.

counting (2 versions) This simple program loops through the range `[1, 2, 3, 4]`, printing "The count is `i`" and then "Done counting" for each number `i`. The **nospace** version

has the "Done counting" `print` statement immediately following "The count is `i`," whereas the **twospaces** version has two blank lines in between. Python is sensitive to horizontal whitespace, but not vertical, so the extra lines do not change the output of the program.

We expected more participants to mistakenly assume that the "Done counting" `print` statement was not part of the loop body in the `twospaces` version. This was the case: 59% of responses in the `twospaces` version contained this error as opposed to only 15% in the `nospace` version (ref = `nospace`, OR = 4.0, $p < 0.0001$). Blank lines, while not usually syntactically relevant, are positively correlated with code readability [2]. We did not find a significant effect of experience on the likelihood of making this mistake, suggesting that experts and novices alike may benefit from an ending delimiter (e.g., an `end` keyword or brackets).

funcall (3 versions) This program prints the product $f(1) * f(0) * f(-1)$ where $f(x) = x + 4$. The **nospace** version has no spaces between the calls to `f`, while the **space** version has a space before and after each `*` (e.g., $f(1) \ * \ f(0) \ * \ f(-1)$). The **vars** version saves the result of each call to `f` in a separate variable, and then prints the product of these variables. Code for `funcall`'s is not included for space reasons.

Most people were able to get the correct answer of 60 in 30 seconds or less. The most common errors (only 7% of responses) were 0, -60, and -80. We hypothesize that these correspond to the following calculation errors: assuming $f(0) = 0$, $f(-1) = -3$, and $f(-1) = -4$. There were no significant effects of version or experience on grade.

initvar (3 versions) The `initvar` program computes the product and sum of variables `a` and `b`. In the `good` version, `a` is initialized to 1, so the product computes $4! = 24$, and `b` is initialized to 0, making the summation 10. In the `onebad` version, $b = 1$, offsetting the summation by 1. In the `bothbad` version, $b = 1$ and $a = 0$, which makes the product 0.

We expected experienced programmers to make more errors due to the close resemblance of code in the `*bad` versions to common operations performed in the `good` version (factorial and summation). Instead, we found a significant negative effect of the `good` version on grade (intercept = 8.67, $\beta = -1.52$, $p < 0.05$), which is likely due to the difficulty of mentally performing 4 factorial. In the `bothbad` version, $a = 0$, allowing participants to short-circuit the multiplication (since a times anything is still zero). The `onebad` version, which also required performing the factorial, had a negative but non-significant effect on grade ($\beta = -0.97$).

order (2 versions) The `order` program prints the values of three functions, $f(x)$, $g(x)$, and $h(x)$. In the `inorder` version, f, g, and h are defined and called in the same order. The `shuffled` version defines them out of order (h, f, g).

We expected programmers to be slower when solving the `shuffled` version, due to an implicit expectation that function definitions and use would follow the same order. When including years of Python experience, we found a significant main effect on RT of the `shuffled` version (intercept = 54.3, $\beta = 21.0$, $p < 0.05$) as well as an interaction between experience and `shuffled` ($\beta = -7.1$, $p < 0.05$). Functions defined out of order had a significant impact on response time, but experience helps counter-act the effect.

overload (3 versions) This program uses the overloaded `+` operator, which serves as addition for integers and concatenation for strings. The `plusmixed` version uses both overloads of the operator (`3 + 7`, `"5" + "3"`), while the `multmixed` version and `strings` version only use `+` for string concatenation (`"5" + "3"`).

We expected programmers in the `plusmixed` version to make the mistake of interpreting `"5" + "3"` as 8 instead of `"53"` more often due to the priming of `+` as addition instead of concatenation. While this error occurred in about 11% of responses across all versions, we did not see a significant grade difference between versions. For response time, a significant interaction between overall programming experience and the `plusmixed` version was found (intercept = 42.5, $\beta = 3.34$, $p < 0.01$). Experienced programmers were slowed down more by the `plusmixed`

version than inexperienced programmers, perhaps due to increased expectations that clustered uses of `+` should correspond to the same operation (addition **or** concatenation).

partition (3 versions) The `partition` program iterates through the ranges $[1, 4]$ (unbalanced) or $[1, 5]$ (balanced), printing out `i low` for $i < 3$ and `i high` for $i > 3$. The `balanced` version outputs two `low` and two `high` lines, whereas the `unbalanced` versions produce two `low` lines and only one `high` line. The `unbalanced_pivot` version calls attention to 3 by assigning it to a variable named `pivot`.

We expected participants in the `unbalanced*` versions to add an additional `high` line because there were four numbers in the list (making it desirable for there to be four lines in the output). While there were a handful of responses like this, the most common error was simply leaving off the numbers on each line (e.g., `low` instead of `1 low`). Programmers seeing the `unbalanced` version were less susceptible to this error (ref = `balanced`, OR = 0.05, $p < 0.05$), though we saw no effect for the `unbalanced_pivot` version. More programming experience also helped participants avoid this kind of mistake across versions (base = 1.66, OR = 0.67, $p < 0.05$). We hypothesize that the `balanced` and `unbalanced_pivot` versions matched a "partition" schema for programmers, making them less likely to pay close attention to the loop body.

rectangle (3 versions) This program computes the areas of two rectangles using an `area` function with x and y scalar variables (`basic` version), (x, y) coordinate pairs (`tuples` version), or a `Rectangle` class (`class` version).

We expected participants seeing the `tuples` and `class` versions to take longer, because these versions contain more complicated structures. Almost everyone gave the correct answer, so there were no significant grade differences between versions. We found a significant RT main effect for the `tuples` version (intercept = 53.5, $\beta = 60.4$, $p < 0.01$), and an interaction between this version and Python experience ($\beta = -34.1$, $p < 0.01$). Programmers in the `tuples` version took longer than those in the `basic` version, but additional Python experience helped reverse this effect. Surprisingly, we did not observe even a marginally significant RT effect for the `classes` version, despite it being the longest program of the three (21 lines vs. 14 and 18).

scope (2 versions) This program applies four functions to a variable named `added`: two `add_1` functions, and two `twice` functions. The `samename` version reused the name `added` for function parameters, while the `diffname` version used `num`. Because Python uses "pass by value" semantics with integers, and because neither of the functions return a value, `added` retains its initial value of 4 throughout the program (instead of being 22). This directly violates one of Soloway's Rules of Discourse [8]: do not include code that will not be used.

We expected participants to mistakenly assume that the value of `added` was changed more often when the parameter names of `add_1` and `twice` were both also named `added` (`samename` version). There was marginally significant evidence for this ($p = 0.09$), but it was not conclusive. Additional Python experience helped reduce the likelihood of answering 22 (base = 1.28, OR = 0.71, $p < 0.05$), but around half of the participants still answered incorrectly.

whitespace (2 versions) The `whitespace` program prints the result of three simple linear calculations. In the `zigzag` version, the code is laid out with one space between every mathematical operation, so that the line endings have a typical zig-zag appearance. The `linedup` version aligns each block of code by its mathematical operators.

We expected there to be a speed difference between the two versions in favor of `linedup`. When designing the experiment, most of our pilot participants agreed that this version was easier to read. The data did not support this claim, but there was a significant effect on the probability of not respecting order of operations. For the `zigzag` version, participants were significantly more likely to incorrectly answer 5, 10, and 15 for the y column (ref = `linedup`, OR = 0.18, $p < 0.05$). These are the answers that would be obtained if a participant executed the multiplications before the additions, contra the established of order of operations of Python

and mathematics more generally. This suggests that when computing the y values, participants in the `zigzag` version did addition before multiplication more often than in `linedup` version. Effects of spacing on the perceived order of arithmetic operations has been studied before [5], and our results indicate that spacing in code layout also has an impact on order of executed operations.

5 Discussion

Experience helps experts in situations where they have reason to monitor for specific kinds of errors, but may hurt in cases for which they have not been trained. For example, our results from the `order` programs show that experience protects programmers from being sensitive to the shuffled order of the functions, because it is often the case in real world programs that functions are defined and used out of order. However, experience leads to more of a tendency to be primed in the `overload` programs because it is unusual to use `+` for addition and then immediately for string concatenation. Real programs tend to have clumps of similar usage of an operator, and programmers learn to be efficient by taking advantage of those frequently occurring repetitions. This same effect can be seen in `between` programs, where experience leads to the expectation that the `common` operation should immediately use the results of the `between` operations. Expectations are sometimes so strong, however, that experience only plays a small role in avoiding errors. Programmers in both versions of the `scope` program strongly expected the `add_1` and `twice` functions to do what their names implied, despite Python's call-by-value semantics for integers and the fact that neither function actually returned a value.

The physical aspects of notation, often considered superficial, can have a profound impact on performance. Programmers were more likely to respect the order of mathematical operations in the `linedup` version of `whitespace`, showing how horizontal space can emphasize the common structure between related calculations. Similarly, the `twospaces` version of `counting` demonstrated that vertical space is more important then indentation to programmers when judging whether or not statements belong to the same loop body. Programmers often group blocks of related statements together using vertical whitespace, but our results indicate that this seemingly superficial space can cause even experienced programmers to internalize the wrong program. Notation can also make a simple program more difficult to read. Programmers took longer to respond to the `tuples` version of `rectangle` despite it having fewer lines than the `class` version. It is not uncommon in Python to use tuples for (x, y) coordinates, but the syntactic "noise" that is present in the `tuples` version for variable names (e.g., `r1_xy_1`) and calculations (e.g., `width = xy_2[0] - xy_1[0]`) likely gave programmers pause when verifying the code's operation.

Future Work During the course of the experiment, Bloomington participants were seated in front of a Tobii X300 eye-tracker. We plan to analyze this eye-tracking data, and correlate it with our findings here. Specifically, we hope to see how code features and experience effect the visual search process and, by proxy, program comprehension.

6 Acknowledgments

Grant R305A1100060 from the Institute of Education Sciences Department of Education and grant 0910218 from the National Science Foundation REESE supported this research.

References

[1] J.M. Burkhardt, F. Détienne, and S. Wiedenbeck. Object-oriented program comprehension: Effect of expertise, task and phase. *Empirical Software Engineering*, 7(2):115–156, 2002.

[2] Raymond Buse and Wes Weimer. Learning a metric for code readability. *Software Engineering, IEEE Transactions on*, 36(4):546–558, 2010.

[3] Huw Talfryn Oakley Davies, Iain Kinloch Crombie, and Manouche Tavakoli. When can odds ratios mislead? *Bmj*, 316(7136):989–991, 1998.

[4] F. Détienne and F. Bott. *Software design–cognitive aspects*. Springer Verlag, 2002.

[5] Robert L. Goldstone, David H. Landy, and Ji Y. Son. The education of perception. *Topics in Cognitive Science*, 2(2):265–284, 2010.

[6] Mark Guzdial. From science to engineering. *Commun. ACM*, 54(2):37–39, February 2011.

[7] T.J. McCabe. A complexity measure. *Software Engineering, IEEE Transactions on*, (4):308–320, 1976.

[8] E. Soloway and K. Ehrlich. Empirical studies of programming knowledge. *IEEE Transactions on Software Engineering*, (5):595–609, 1984.

A Appendix

A.1 Programs

A.1.1 between - functions

```
def between(numbers, low, high):
    winners = []
    for num in numbers:
        if (low < num) and (num < high):
            winners.append(num)
    return winners

def common(list1, list2):
    winners = []
    for item1 in list1:
        if item1 in list2:
            winners.append(item1)
    return winners

x = [2, 8, 7, 9, -5, 0, 2]
x_btwn = between(x, 2, 10)
print x_btwn

y = [1, -3, 10, 0, 8, 9, 1]
y_btwn = between(y, -2, 9)
print y_btwn

xy_common = common(x, y)
print xy_common
```

Output:

```
[8, 7, 9]
[1, 0, 8, 1]
[8, 9, 0]
```

A.1.2 between - inline

```
x = [2, 8, 7, 9, -5, 0, 2]
x_between = []
for x_i in x:
    if (2 < x_i) and (x_i < 10):
        x_between.append(x_i)
print x_between

y = [1, -3, 10, 0, 8, 9, 1]
y_between = []
for y_i in y:
    if (-2 < y_i) and (y_i < 9):
        y_between.append(y_i)
```

```
print y_between

xy_common = []
for x_i in x:
    if x_i in y:
        xy_common.append(x_i)
print xy_common
```

Output:

```
[8, 7, 9]
[1, 0, 8, 1]
[8, 9, 0]
```

A.1.3 counting - nospace

```
for i in [1, 2, 3, 4]:
    print "The count is", i
    print "Done counting"
```

Output:

```
The count is 1
Done counting
The count is 2
Done counting
The count is 3
Done counting
The count is 4
Done counting
```

A.1.4 counting - twospaces

```
for i in [1, 2, 3, 4]:
    print "The count is", i


    print "Done counting"
```

Output:

```
The count is 1
Done counting
The count is 2
Done counting
The count is 3
Done counting
The count is 4
Done counting
```

A.1.5 funcall - nospace

```
def f(x):
    return x + 4

print f(1)*f(0)*f(-1)
```

Output:

```
60
```

A.1.6 funcall - space

```
def f(x):
    return x + 4

print f(1) * f(0) * f(-1)
```

Output:

```
60
```

A.1.7 funcall - vars

```
def f(x):
    return x + 4

x = f(1)
y = f(0)
z = f(-1)
print x * y * z
```

Output:

```
60
```

A.1.8 initvar - bothbad

```
a = 0
for i in [1, 2, 3, 4]:
    a = a * i
print a

b = 1
for i in [1, 2, 3, 4]:
    b = b + i
print b
```

Output:

```
0
11
```

A.1.9 initvar - good

```
a = 1
for i in [1, 2, 3, 4]:
    a = a * i
print a

b = 0
for i in [1, 2, 3, 4]:
    b = b + i
print b
```

Output:

```
24
10
```

A.1.10 initvar - onebad

```
a = 1
for i in [1, 2, 3, 4]:
    a = a * i
print a

b = 1
for i in [1, 2, 3, 4]:
    b = b + i
print b
```

Output:

```
24
11
```

A.1.11 order - inorder

```
def f(x):
    return x + 4

def g(x):
    return x * 2

def h(x):
    return f(x) + g(x)

x = 1
a = f(x)
b = g(x)
c = h(x)
print a, b, c
```

Output:

```
5  2  7
```

A.1.12 order - shuffled

```
def h(x):
    return f(x) + g(x)

def f(x):
    return x + 4

def g(x):
    return x * 2

x = 1
a = f(x)
b = g(x)
c = h(x)
print a, b, c
```

Output:

```
5  2  7
```

A.1.13 overload - multmixed

```
a = 4
b = 3
print a * b

c = 7
d = 2
print c * d

e = "5"
f = "3"
print e + f
```

Output:

```
12
14
53
```

A.1.14 overload - plusmixed

```
a = 4
b = 3
print a + b
```

```
4
5  c = 7
6  d = 2
7  print c + d
8
9  e = "5"
10 f = "3"
11 print e + f
```

Output:

```
7
9
53
```

A.1.15 overload - strings

```
1  a = "hi"
2  b = "bye"
3  print a + b
4
5  c = "street"
6  d = "penny"
7  print c + d
8
9  e = "5"
10 f = "3"
11 print e + f
```

Output:

```
hibye
streetpenny
53
```

A.1.16 partition - balanced

```
1  for i in [1, 2, 3, 4, 5]:
2      if (i < 3):
3          print i, "low"
4      if (i > 3):
5          print i, "high"
```

Output:

```
1 low
2 low
4 high
5 high
```

A.1.17 partition - unbalanced

```
for i in [1, 2, 3, 4]:
    if (i < 3):
        print i, "low"
    if (i > 3):
        print i, "high"
```

Output:

```
1 low
2 low
4 high
```

A.1.18 partition - unbalanced_pivot

```
pivot = 3
for i in [1, 2, 3, 4]:
    if (i < pivot):
        print i, "low"
    if (i > pivot):
        print i, "high"
```

Output:

```
1 low
2 low
4 high
```

A.1.19 rectangle - basic

```
def area(x1, y1, x2, y2):
    width = x2 - x1
    height = y2 - y1
    return width * height

r1_x1 = 0
r1_y1 = 0
r1_x2 = 10
r1_y2 = 10
r1_area = area(r1_x1, r1_y1, r1_x2, r1_y2)
print r1_area

r2_x1 = 5
r2_y1 = 5
r2_x2 = 10
r2_y2 = 10
r2_area = area(r2_x1, r2_y1, r2_x2, r2_y2)
print r2_area
```

Output:

```
100
25
```

A.1.20 rectangle - class

```
class Rectangle:
    def __init__(self, x1, y1, x2, y2):
        self.x1 = x1
        self.y1 = y1
        self.x2 = x2
        self.y2 = y2

    def width(self):
        return self.x2 - self.x1

    def height(self):
        return self.y2 - self.y1

    def area(self):
        return self.width() * self.height()

rect1 = Rectangle(0, 0, 10, 10)
print rect1.area()

rect2 = Rectangle(5, 5, 10, 10)
print rect2.area()
```

Output:

```
100
25
```

A.1.21 rectangle - tuples

```
def area(xy_1, xy_2):
    width = xy_2[0] - xy_1[0]
    height = xy_2[1] - xy_1[1]
    return width * height

r1_xy_1 = (0, 0)
r1_xy_2 = (10, 10)
r1_area = area(r1_xy_1, r1_xy_2)
print r1_area

r2_xy_1 = (5, 5)
r2_xy_2 = (10, 10)
r2_area = area(r2_xy_1, r2_xy_2)
print r2_area
```

Output:

```
100
25
```

A.1.22 scope - diffname

```
def add_1(num):
    num = num + 1

def twice(num):
    num = num * 2

added = 4
add_1(added)
twice(added)
add_1(added)
twice(added)
print added
```

Output:

```
4
```

A.1.23 scope - samename

```
def add_1(added):
    added = added + 1

def twice(added):
    added = added * 2

added = 4
add_1(added)
twice(added)
add_1(added)
twice(added)
print added
```

Output:

```
4
```

A.1.24 whitespace - linedup

```
intercept = 1
slope     = 5

x_base  = 0
x_other = x_base + 1
```

```
x_end   = x_base + x_other + 1

y_base  = slope * x_base  + intercept
y_other = slope * x_other + intercept
y_end   = slope * x_end   + intercept

print x_base,  y_base
print x_other, y_other
print x_end,   y_end
```

Output:

```
0 1
1 6
2 11
```

A.1.25 whitespace - zigzag

```
intercept = 1
slope = 5

x_base = 0
x_other = x_base + 1
x_end = x_base + x_other + 1

y_base = slope * x_base + intercept
y_other = slope * x_other + intercept
y_end = slope * x_end + intercept

print x_base, y_base
print x_other, y_other
print x_end, y_end
```

Output:

```
0 1
1 6
2 11
```

A.2 Tables

Table 1: Results by program version. (*) = log. regression reference, LOC = lines of code, CC = cyclomatic complexity, RT = response time. Main effects listed for version and experience. CE = prob. of common error, GR = grade, * = significance.

Type	Version	LOC	CC	Avg. Grade	Mean RT (s)	Effects (ver)	(exp)	(ver × exp)
between	functions (*)	24	7	4.7	142.8		CE ↑ *	
	inline	19	7	5.8	151.5			
counting	nospace (*)	3	2	8.8	66.6			
	twospaces	5	2	5.9	55.6	CE ↑ ***		
funcall	nospace (*)	4	2	9.1	38.6			
	space	4	2	8.8	35.9			
	vars	7	2	9.8	36.9			
initvar	bothbad (*)	9	3	8.7	63.2			
	good	9	3	7.1	66.0	GR ↓ *		
	onebad	9	3	7.7	61.6			
order	inorder (*)	14	4	8.7	61.4			
	shuffled	14	4	9.1	68.1	RT ↑ *		RT ↓ *
overload	multmixed (*)	11	1	8.9	37.3			
	plusmixed	11	1	8.7	41.6			RT ↑ **
	strings	11	1	8.5	39.6			
partition	balanced (*)	5	4	6.9	45.9			
	unbalanced	5	4	8.0	41.6	CE ↓ *		
	unbalanced_pivot	6	4	8.1	39.6			
rectangle	basic (*)	18	2	9.7	76.5			
	class	21	5	9.4	72.5			
	tuples	14	2	9.5	80.1	RT ↑ **		RT ↓ **
scope	diffname (*)	12	3	7.2	58.0		CE ↓ *	
	samename	12	3	6.7	57.9			
whitespace	linedup (*)	14	1	8.7	111.7			
	zigzag	14	1	8.5	108.4	CE ↑ *		